

**This Page Is Inserted by IFW Operations
and is not a part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
Please do not report the images to the
Image Problem Mailbox.

This Page Blank (uspto)

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

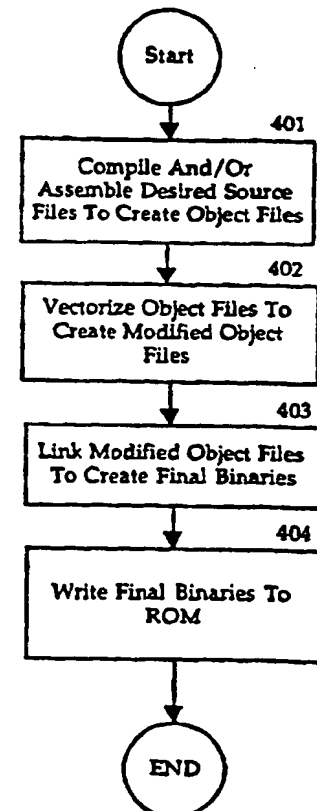
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁵ : G06F 9/44		A1	(11) International Publication Number: WO 94/27219
			(43) International Publication Date: 24 November 1994 (24.11.94)
(21) International Application Number: PCT/US94/04995			(81) Designated States: AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, ES, FI, GB, GE, HU, JP, KG, KP, KR, KZ, LK, LU, LV, MD, MG, MN, MW, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SI, SK, TJ, TT, UA, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 6 May 1994 (06.05.94)			
(30) Priority Data: 08/058,877 6 May 1993 (06.05.93) US			
(71) Applicant: APPLE COMPUTER, INC. [US/US]; 20525 Marina Avenue, Cupertino, CA 95014 (US).			
(72) Inventors: WETMORE, Russ; 11628 Birch Springs Court, Cupertino, CA 95014 (US). NGUYEN, Philip; 475 Alto Drive, Boulder Creek, CA 95006 (US). BATISTA, Ricardo; 5921 Mivaga Avenue, San Jose, CA 95123 (US).			
(74) Agents: SCHELLER, James, C. et al.; Blakely, Sokoloff, Taylor and Zisman, 12400 Wilshire Boulevard, 7th floor, Los Angeles, CA 90025 (US).			Published With international search report.

(54) Title: METHOD AND APPARATUS FOR PATCHING CODE RESIDING ON A READ ONLY MOEMORY DEVICE

(57) Abstract

A method and apparatus for generating patching resources in an information processing system having operating instructions on a Read Only Memory Device. The present invention simplifies the patch generation and installation processes. A patch resource is generated and used by a patch installation process. Patch resources are generated for each ROM version by comparing previous ROM versions to the new ROM version. A patch resource is comprised of a plurality of entries, each of which defines a vector table address, an offset into the vector table and the routine to be inserted. By comparing routines between the ROM versions, routines which are different or new are identified. These routines will become patch resource entries. For patch installation, the ROM version number for the installed ROM is determined; the proper patching resource is retrieved, and the patch resource entries cause the patches to be installed. Patch installation is performed by the steps of modifying vector tables to include the addresses for the new routines.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

-1-

METHOD AND APPARATUS FOR PATCHING CODE RESIDING ON A READ ONLY MEMORY DEVICE

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to the field of computer operating systems and the underlying code structure thereof, in particular to applying fixes to code residing on static storage such as a Read Only Memory device.

2. Description of the Related Art

It is well known that a computer system relies on operating software, i.e. an operating system, to enable basic computer functionality. For example, an operating system allows a user to store and retrieve files on a storage medium. Various approaches are used to provide the operating system software as part of the computer system. One approach utilized in IBM compatible computer systems is to provide a Basic Input Output System (BIOS) on a Read Only Memory (ROM) device. The BIOS contains the instructions for interaction between the various components of the computer system. The remainder of the Operating system functionality is loaded in Random Access Memory (RAM). In such implementations, the vast majority of the operating system functionality is loaded into the RAM. Other aspects of the operating environment namely the user interface tools, also would exist on RAM. This approach has the drawback of utilizing RAM which could otherwise be used for application programs.

-2-

An alternative approach is to provide as much operating system functionality into ROM as possible. This has the desired effect of freeing up RAM for application programs. This approach is used for the operating system for the Apple® Macintosh® family of computers, available from Apple Computer, Inc. of Cupertino, California. The organization of the Macintosh operating software between ROM and RAM as well as the Macintosh environment in general, is discussed in the publication entitled "Inside Macintosh Volume I", available from Addison-Wesley Publishing company.

The portion of the Macintosh operating environment that resides in ROM is comprised of two parts; the operating system and the user interface toolboxes. The operating system portion provides traditional operating system functionality. The toolboxes provides a standardized set of tools for application development. Examples of toolboxes include the QuickDraw Manager (for drawing figures on a display) Sound Manager and Resource Manger. The use of such toolboxes would be well known to one having familiarity with developing applications for the Apple Macintosh family of computers.

In the Macintosh environment routines based in ROM are typically accessed using what is known as the A-Trap dispatching mechanism. The A-Trap dispatching mechanism is described in the publication "Programmer's Guide To MPW Volume 1", Mark Andrews, available from Addison-Wesley publications (MPW is an acronym for Macintosh

-3-

Programmer's Workshop). The A-Trap dispatching mechanism allows for the calling of the ROM based routines symbolically through the trap dispatcher, rather than by absolute ROM address.

One problem with storing code in ROM is that it is static and cannot be fixed (absent physically replacing and re-writing the ROM). Accordingly, adding functionality or fixing "bugs" found in the operating system ROM code is very tricky. To fix a bug or add functionality, one must either patch the vectors maintained by the A-Trap dispatching mechanism, or patch the private vectors maintained by some of the tool box managers. "Patch" is a term of art which refers to new code introduced to fix prior code or to add functionality. A ROM vector causes a jump to a location in RAM where the patch code may reside. However, because there are a limited number of such vectors, most of the code is called directly and cannot be easily patched.

To patch non-vectorized code, one must be very creative. In some cases, all clients of the offending code can be patched. Clients in this context refers to code that calls or receives data from the offending code. In other cases a routine called by the offending routine may be patched to fix what the calling routine did wrong. This is called a "come from" patch and it usually identifies the caller by comparing the return address with a known absolute address. When small patches are made to large routines, it is common practice to call the existing code in ROM to save memory. Usually

-4-

this is done by jumping to the absolute address in the ROM. In doing so, the absolute address in ROM becomes hard coded into the patch.

Because of these absolute addresses hard-coded into the patches, the ROM is very difficult to maintain. Much care must be taken to assure any changes or additions to the ROM will not change the addresses of the existing code. This has the undesirable effect of making the ROM based code non-relocatable (because of code reliance on absolute addresses). This becomes even more difficult as more operating system code is written in high level languages.

Despite such obstacles, it is desirable to place operating system functionality in ROM because it reduces the amount needed for RAM. Consequently, this frees RAM resources to be used for application software programs. Another advantage is that it is easier to protect ROM based code from unauthorized copying.

Moreover, as application software becomes integrated into base functionality of computer system, it is likely the application software itself will become ROM based. As more functionality is placed in ROM, the foregoing maintenance difficulties are compounded.

Thus, it is an object of the present invention to provide a mechanism for generating code that will reside in ROM so that patches or additional function may be added with greater ease. It is a further object of the present invention to simplify the patch installation process.

-5-

SUMMARY

The present invention is directed towards a method and apparatus for generating patches for static operating instructions executing on a computer system. The present invention is particularly useful in the generation of programs that are stored on a static storage device such as a Read Only Memory (ROM) device. Since the code on ROM cannot be fixed without removing the ROM from the computer system, patches must be implemented by vectors into the ROM code that resides on system Random Access Memory (RAM). As the introduction of new operating software must take into account various previous versions of the operating software, the task of patch generation can be arduous.

The present invention simplifies the patch generation and installation processes. Patch resources embody the different routines between the different routines in the various ROM versions. Patch resources are generated for each ROM version by comparing the previous ROM versions to the new ROM version. A patch resource defines a vector table address, an offset into the vector table and the routine to be inserted. By comparing routines of the ROM versions, routines which are different or new are identified. These routines will then be used to create patch resource entries. A patch resource entry effectively provides for replacement of a routine, adding routines or the addition of new functions.

For patch installation, the ROM version number for the installed ROM is determined; the proper patching resource is retrieved, and the patch

-6-

resource entries cause the patches to be installed. Patch installation is performed through modification of existing vector table entries (replacing a routine), adding new vector table entries (adding new routines to an existing function) or adding a new vector table (adding new routines for new functions).

-7-

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram illustrating the basic components of a computer system as may be utilized by the currently preferred embodiment of the present invention.

Figure 2 illustrates the general organization of Random Access Memory (RAM) of the currently preferred embodiment of the present invention.

Figure 3 illustrates the implementation of an indirect vector as may be utilized in the currently preferred embodiment of the present invention.

Figure 4 is a flowchart which describes the steps of creating a vectorized ROM as may be performed in the currently preferred embodiment of the present invention.

Figure 5 is a flowchart which describes the specific steps for vectorizing an object file as may be performed in the currently preferred embodiment of the present invention.

Figure 6 is a chart illustrating a Vector Table Source File as may be utilized by the currently preferred embodiment of the present invention.

-8-

Figure 7a is a flowchart describing the steps for generating vector patch resources, as may be performed by the currently preferred embodiment of the present invention.

Figure 7b is a flowchart describing the steps for vector table updating that is performed during patching, as may be performed by the currently preferred embodiment of the present invention.

-9-

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

This specification is related to, and hereby fully incorporates by reference, serial no. _____, entitled "Method And Apparatus For Vectorizing The Contents Of A Read Only Memory Device Without Modifying Underlying Source Code", filed on _____, and assigned to the same assignee, Apple Computer, Inc..

A method and apparatus for patching coded instructions residing in a Read Only Memory (ROM) device, in a computer system is described. In the following description, numerous specific details are set forth such as coding examples, in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known circuits, control logic and coding techniques have not been shown in detail in order not to unnecessarily obscure the present invention.

The following description will include various code examples of assembly language instructions of the Motorola 680X0 family of microprocessors. Further, various references are made to the structure of the operating environment of the Apple Macintosh. In both instances, one familiar with programming in the Macintosh environment would be familiar with such references and related concepts.

-10-

Overview of the Computer System
of the Preferred Embodiment

The computer system of the preferred embodiment is described with reference to Figure 1. The present invention is preferably implemented on a general purpose microcomputer in which a significant amount of operating or application software resides on a static memory device, such as one of the members of the Apple® Macintosh® family of computers. In any event, a computer system as may be utilized by the preferred embodiment generally comprises a bus or other communication means 101 for communicating information, a processing means 102 coupled with said bus 101 for processing information, a random access memory (RAM) or other storage device 103 (commonly referred to as a main memory) coupled with said bus 101 for storing information and instructions for said processor 102, a read only memory (ROM) or other static storage device 104 coupled with said bus 101 for storing static information and instructions for said processor 102, a data storage device 105, such as a magnetic disk and disk drive, coupled with said bus 101 for storing information and instructions, an alphanumeric input device 106 including alphanumeric and other keys coupled to said bus 101 for communicating information and command selections to said processor 102, a cursor control device 107, such as a mouse, track-ball, cursor control keys, etc., coupled to said bus 101 for communicating information and command selections to said processor 102 and for controlling cursor movement, and a display device 108 coupled to bus 101 for displaying textual, graphical or video output. Additionally, it is useful if the system includes a hardcopy device 109, such

-11-

as a printer, for providing permanent copies of information. The hardcopy device 109 is coupled with the processor 102 through bus 101.

The computer system of Figure 1 may be both a system which generates the vectorized ROM code or a system which utilizes the vectorized ROM code, or both. In a computer system having the vectorized code of the present invention, it is the ROM 104 that will contain the static code. Non-static code and vector tables will reside in the RAM 103. Further, during the generation of such vectorized code, a processor on a computer system, such as the processor 102 of Figure 1, will perform processing function means. For example, the vectorizing and linking of object files would be performed by the processor 102.

The organization of RAM 103 as found on the preferred Apple Macintosh System is illustrated in Figure 2. For the purposes of this description, consider the RAM as organized into two areas: system heap 201 and application space 202. The system heap area 201 contains all the operating system code and utilities. As will be described in more detail below, it is in this system heap area 201 that the vector tables are loaded and stored. The application space 202 is the available memory for applications programs.

OVERVIEW OF ROM VECTORIZATION

As noted above with respect to the prior art, there are numerous benefits of having ROM based code. However, as code in ROM is static,

-12-

patching to apply fixes or add functionality is tricky and requires patching to an existing vector. Although vectors are used on the A-Trap mechanism and locally by some toolbox managers, the number of vectors is far too few for maintaining a dynamic and robust operating environment.

Consequently the present invention provides a means for generalizing the introduction of vectors into ROM based code. It should be noted that the fact that this code is to reside on ROM is not meant to limit the scope of the present invention. Any system that utilizes "static" code on, for example "FLASH" memory magnetic or optical disk media or other storage devices, could be utilized without causing departure from the spirit and scope of the present invention.

Another benefit of this technique is that the programmer need not be constrained or concerned about the absolute physical addresses in the ROM. All the programmer has to be concerned about is generating the source code. Moreover as there is no change to the source code, the use of vectors is transparent to the programmer. This will greatly simplify ROM maintenance and allow products to be developed and introduced into the marketplace at a much quicker pace.

ROM vectorization is the process by which static program code that is to be installed in ROM is modified to create external references to a vector table in RAM. The entries in the vector table contain pointers to the location of the various code that will be executed. Generally, when the ROM code is vectorized, entry points for external routines are replaced by a

-13-

reference to a table and an offset into the table. The corresponding table entry will then point to the location of the routine. So as the ROM code is executing, upon encountering a reference to an external routine, e.g. a subroutine or function call, the actual entry point will reference the vector table and the corresponding entry in the vector table will point to (e.g. have the address of, or be a JMP instruction to) the actual code to be executed.

The effect of vectorization is described with reference to Figure 3. Referring to Figure 3, illustrated is ROM based code 301 and RAM code 302. The only thing illustrated here in the RAM code 302 is the vector table 306. In any event, ROM code 301 will contain a reference to an entry point 303. The entry point may be a sub-routine, function, macro or a jump to a label somewhere else in the ROM code 301. Note that various other linkage code to allow return after the sub-routine or function is not illustrated but is assumed to exist. The reference to entry point 303 will effectively point to entry point 304. Without vectorization, the executable code would be immediately following the entry point 304. With vectorization, the location of entry point 304 has been modified to be a reference to a table pointer 305 that resides in RAM. The table pointer 305 is the vector in this example. The table pointer 305 will point to vector table 306 which resides in RAM code 302 (specifically the system heap area).

As noted above, the reference to the table would in most cases include an offset into the vector table 306. Assuming the offset, the

-14-

entry 307 will contain a pointer to the location where the code to be executed would reside. While the entry 307 may simply point back into the ROM, in the instance of a patch, the entry 307 may contain a pointer to an updated routine located somewhere in the RAM 302.

In the currently preferred embodiment, the foregoing example describes an "Indirect Vector" type (i.e. the table pointer 305). The currently preferred embodiment includes two other vector types; a Direct Vector and an Indirect Jump Vector. A Direct Vector is longword in low memory that contains the address of the routine. A Direct Vector is used when execution speed is the paramount concern. The Indirect Vector Jump is similar to the Indirect Vector but differs by using a Jump table in place of a vector table. The Indirect Vector Jump is used in situations when a scratch register isn't available and the code is in cache.

Vectorization facilitates maintenance of ROM based code by removing the need to rely on actual physical addresses in ROM when fixing "bugs". The term "bug" is a term well known to those skilled in the art and in the context of software design, refers to the existence of logic or other errors which causes a software program to malfunction. Basically, through the vectorization of the present invention, more entry points into the ROM are created, thus providing more locations at which the ROM may be accessed and code fixed. Further, it eliminates the need to hard code absolute addresses into the patched code. In other words, the ROM code is

-15-

modularized to a greater extent so that respective modules may be more easily replaced.

Organization of Vector Tables

In the Macintosh environment, the various application development tools are organized into toolboxes. Control of a toolbox is performed by it's "toolbox manager". In the currently preferred embodiment each "toolbox manager" will have it's own vector table. For example, in the Apple Macintosh environment, the Window manager, Dialog manager and QuickDraw manager all have their own vector tables. By arranging the vector routines in groups the tables in system software may be easily expanded. Also, complete tables can be replaced with new one when a tool box manager is rewritten. Each vector table is accessed through a pointer stored in low memory.

Another advantage of organization into vector groups, is that internal ROM code can make use of the vector tables directly. New ROM source code could be developed taking advantage of the vector groups assigned so far. Using this technique a vector table does not need to have a predefined size.

By keeping the vectors in their own respective groups, vector table initialization may occur independently (as will be described in greater detail below, vector tables must be initialized before use). Vectorized routines must have their vector table entry initialized. This is accomplished by

-16-

creating a small routine for each vector group which takes care of the vector table initialization. As will be described below, the initialization routine is created during the vectorization process.

ROM Vectorization

The manner in which the code in a ROM is vectorized is illustrated by the steps in the flowchart of Figure 4. First, the source files are compiled (in the case of a high level language) or assembled (in the case of assembler language source) to create object files, step 401. The object files are then vectorized to create vectorized object files, step 402. It is significant that only the object files are modified. The source files are not touched. Object files contain a series of defined records, each one containing specific items such as the object code for a routine, the name of a routine, external reference from one routine to another, or comments. In object files the references to other routines have not been resolved. Therefore object files are an ideal place to alter the code without modifying the source code files. The steps for vectorization are described in more detail below with respect to Figure 5.

The object files are then linked together to create the final binary values which will be written to ROM, step 403. This is performed through a traditional linkage editing step. Finally, after the object files have been "linked" together to create the final binaries, the ROM image is created, step 404.

-17-

Figure 5 is a flowchart illustrating the steps for vectorizing an object file. Referring to Figure 5, the entry points of the object file are first identified to create a vector table source file, step 501. An entry point may be the name of a routine or a label in the file. Generally, an entry point is merely a location in the code which may be entered via a symbolic reference. It is these entry points which become the code access points which are vectorized. The vector table source file is a set of assembly language instructions. The vector table source file is described in greater detail below. Next, the vector table is assembled in order to create a vector table object file, step 502. For each entry in the vector table source file there is a corresponding module in a vector object file. Each of these modules has one entry point with the vector's name and one content record containing the glue code used to patch the original routine. An example of modules in the vector table object file are illustrated in Table 1.

MaxBlock	Proc	Export
	jmp	([\$0584])
	EndProc	
SwapZone	Proc	Export
	move.1	\$2050,a0
	move.1	\$08(a0), a0
	jmp	(a0)
	EndProc	
SendBit	Proc	Export
	move.1	\$2060,a0

-18-

	jmp	\$06(a0)
	EndProc	

Table 1 Vector Table Object File

Referring to Table 1, three (3) entry points, **MaxBlock**, **SwapZone** and **SendBit** are illustrated. Each of the three entry points includes the "vector code" for accessing the routine. The vector code presented, which is written in the Motorola 680X0 Assembler language, is exemplary. It would be apparent to one skilled in the art that the functionality performed by this code can be implemented using different instruction formats or a different Assembler Language (one supported by the processor of the underlying computer system).

The entry **MaxBlock** is an example of a Direct Vector. The instruction **jmp ([\$0584])** will cause a direct jump to the location of the routine. Here the address of the desired routine is contained in the memory location address **\$0584**. Note that the term jump as used here refers to program execution to continue at the address jumped to.

The entry **SwapZone** is an example of an Indirect Vector. Here the instruction **move.l \$2050,a0** moves the contents of the location **\$2050** (the address for the vector table) into register **a0**. The instruction **move.l \$08(a0),a0** causes the offset **\$08** to be added to the contents of the register **a0**. At this point the register **a0** contains the address for the vector table entry of the desired routine. The **jmp(a0)**

-19-

instruction causes a jump to the contents of the vector table entry, which is the entry address of the desired routine.

The entry **SendBit** is an example of an Indirect Jump Vector. The instruction **move.1 \$2060,a0** causes the contents of address **\$2060** to be moved into the register **a0**. The instruction **jmp \$06(a0)** causes a jump to the address that is offset by **\$06** from the contents of the register **a0**. The address jumped to will contain another jump instruction to the entry address of the desired routine.

Once the vector table object file is created, the vectorized object file is created by replacing the entry point references with the appropriate vector code. Referring back to Figure 5, a symbol table containing the vector names and the vector code is created, step 503. Each of the object files are then processed by comparing entry point names to the names in the symbol table, step 504. If a match is found, the entry point name in the object file is changed and the vector code is inserted in the object file, step 505. This will effectively cause the linkage editor to reference the vector code for all calls made to the original entry point. If no match is found an error/warning message is issued, step 506. After all the object files are vectorized, the vector table initialization code is generated from the vector source table (using a different set of macros than that used to create the vector table object file), step 507.

-20-

For maintenance purposes, each vectorized routine can be a version number so that updates and additions to the routine can be made. Updates and additions to routines is described in greater detail below.

-21-

Table 2 is an example of a routine which has been vectorized.

BEFORE VECTORIZATION		
MaxBlock	Proc	Export
	link	a6,#04
	unlk	a 6
	rts	
	EndProc	
AFTER VECTORIZATION		
__v__MaxBlock	Proc	Export
	link	a6,#04
	unlk	a 6
	rts	
MaxBlock	jmp	([\$0584])
	EndProc	

Table 2 Vectorization Example

Referring to Table 2, before vectorization a routine **MaxBlock** performs the instructions between the code **PROC Export** and **ENDPROC**. In this example, **MaxBlock** is vectorized to contain a direct vector.

After vectorization, the entry name label has been changed to **__v__MaxBlock**. The vector code with the original label **MaxBlock** is then appended to the original code sequence **__v__MaxBlock**. Here the

-22-

label **MaxBlock** has the **jmp ([**\$0584**])** instruction. Absent any patches, the location **\$0584** will contain the address to the label **__v__MaxBlock**.

Starting a System with a Vectorized ROM

As noted above, a vectorized ROM does require that vector initialization code be called before a vector is used. ROM source code does not need to be modified except for the calling of the corresponding initialization routine. The initialization routine must be called from within the ROM before any vectorized routine can be used. This is done automatically at boot (system start-up) time.

Each toolbox manager in ROM has its own vector table pointer in low memory. In the currently preferred embodiment these memory locations have been pre-assigned and are stored in a text file where all of the vector information is kept. When a vector initialization routine is called, it initializes the vector table pointers so that they point to the right location in ROM. Generally, the vector initialize routine will allocate memory in the system heap for the vector table, put the vector table address in the given low memory location, and set-up the vector table with the routine's addresses.

In the currently preferred embodiment, the initialization routine is called three (3) times during the boot process. It is called first when it is determined that there is memory available, then again after the Memory

-23-

Management Unit for the computer system has been setup and finally after the system heap has been created and the system memory initialized to FF. After this third call, the vector pointers are stable and patching can take place. Patching is described in more detail below. However, it would be apparent to one skilled in the art that initialization could be implemented so that it occurs at other times during the system start-up process. Such implementations would not depart from the spirit and scope of the present invention.

Vector Directory

In the currently preferred embodiment the ROM will also contain a directory that describes all the various vectors. This is provided to facilitate the use of program debuggers. Generally, the address of the vector directory will be maintained at a predetermined location in ROM. The information concerning vectors is organized by vector type. Consequently, the vector directory contains pointers to various vector information tables. The vector information tables will contain a count of the vectors in the table, as well as pointers to the name of the corresponding entry point and the original code in ROM.

Vector Table Source File

As is apparent from the foregoing description, in the currently preferred embodiment vectors are implemented using assembly language source files. This provides flexibility at little cost and allowing the performance of conditional compilation. Each entry in the vector table

-24-

source file will have a format depending upon the vector type. An example of the Vector Table Source file is provided in Figure 6. The following is a description of the various fields in the vector source file.

Vector Name is the case sensitive name of the routine to be vectorized.

Vector Type is a macro that specifies which type of vector to apply to the particular routine.

Vector Table Handle is the address where the pointer to the particular vector table can be found at runtime. For direct vectors, the field makes little sense and should be zero. It would be possible to allocate a vector table pointer for direct vectors that could point to address zero or to the base of the appropriate trap dispatch table. This would allow all the vector utility code to work the same regardless of the vector type.

Vector Table Offset is the offset into the vector table. For direct vectors, this is the absolute address of the vector itself.

Dispatch Register identifies a register that can be used to optimize the routine dispatch. If this field is zero or omitted, no optimization will take place.

Runtime Conditions This field can be used to select which code to install at runtime. The constants given must be compatible with the test for macro.

As described above, during the vectorization process, the vector table source file is compiled to produce an object file used by the

-25-

vectorization tool. It is then recompiled using an alternate set of macros to produce the code that initializes the vector table.

Patching A Vectorized ROM

As the purpose of vectorizing the code is to facilitate fixing bugs or adding functionality, it is now useful to describe how it is done. The term patching is used to describe the process for creating and installing patches to the ROM that add functionality or fix bugs. In this particular instance, we are talking about patches applied to ROM vectors

One difficulty in making patches results from the need to support prior versions of the ROM and the ROM code with each released ROM the vectorized routine will have a version number. If a bug is discovered or a new function is added to the routine, then the new routine will have a higher version number when distributed with the new system disk. A ROM maintenance data base will keep track of all different versions of all the vectorized routines as well as which version belongs to what ROM and is smart enough to include the right version for each system release.

The patch mechanism of the currently preferred embodiment creates a vector patch resource to contain the new vectorized routines. A resource in the Macintosh environment refers to a static block of data that may be created, stored and manipulated independently from the program code. The vector patch resource is distributed on the system disk.

-26-

The system disk contains the portion of the operating system environment that does not reside on ROM. In the currently preferred embodiment. The system disk is used to create the operating environment for a computer.

-27-

In the currently preferred embodiment, an entry in the vector patch resource has the format illustrated in Table 3.

<u>FIELD</u>	<u>USE</u>
VectorTable Pointer	Pointer to Vector Table in Low Memory
VectorTable Entry	Offset into Vector Table For Entry For The Routine
Size of vectorized Routine	Specific Size of Routine in Bytes
Vectorized Routine	The New Code to Be Inserted

Table 3 Vector Patch Resource Entry Format

Referring to Table 3, the **VectorTable Pointer** and **VectorTable Entry** are used to identify the Vector Table and the entry for the routine in the vector table, respectively, corresponding to the code that is to be inserted. The size of the Vectorized Routine precedes the actual Vectorized code.

The vector patch resource will typically contain numerous entries corresponding to the number of patches or the added functionality being provided. Note that there will be a vector patch resource for each version of the ROM that supports vectorized routines. Each vector patch resource will have an identifier corresponding to the ROM versions on which it should be loaded. So during the installation process the proper vector patch resource must be identified.

Vector patch resources are created when the operating system is updated and installed when the operating system is "built". The operating

-28-

system is "built" whenever a user wishes to update their computer operating system software to a later release or version level. In the currently preferred embodiment of the present invention, a tool termed ROMPatch, is provided which automatically creates the vector patch resources. ROMPatch compares the object files of two versions of the vectorized ROM code to identify routines which are different or new. In the currently preferred embodiment, routines which are different is accomplished via a Cyclical Redundancy Check (CRC) operation. However, other techniques, e.g. assigning each routine a version number and simply comparing these version numbers, may be utilized without departure from the spirit and scope of the present invention. In any event, when all the patched routines are found, the vector patch resource is generated.

The operation of the ROMPatch tool is further described with reference to Figure 7a. Referring to Figure 7a, the version information of routines of a first (previous) ROM version to a second (new) ROM version, step 701. As described above, a CRC operation may be performed between corresponding routines to determine if it has been changed. For each routine that is identified as a replacement routine, i.e. a new routine that will replace an existing routine, a vector patch resource entry for a replacement patch is created, step 702. A routine may be identified as a replacement routine by determining that the routine exists in both versions. For each routine that is identified as a new routine for an existing function, a vector resource patch entry for adding a routine to an existing function is

-29-

generated, step 703. Here, a function will have its own vector table. So this will involve adding an entry to an existing vector table. Finally, for each new function routine, a vector patch resource entry for adding new function is generated, step 704. This will involve causing a new vector table to be created as well as the offset for the entries to be loaded into the vector table. The information for creation of new tables would come from the new ROM version object file.

To perform the patching a NewVector loader is included with the system disk and its sole purpose is to update and add vectorized routines. At boot time, the vector patch resource from the system files are loaded and only the vector patch resources with ID equal to or greater than the version of the ROM will be loaded. The operation of the NewVector loader is described with respect to the flowchart in Figure 7b. First, the vector patch resource corresponding to the ROM version of the system being updated is identified, step 721. The remaining steps are performed for each entry in the vector patch resource that has been identified. When the vector patch resource is received, the entry must be identified as a replacement of an old routine, new functionality or a new routine, step 722. A determination is made if it is the replacement of an old routine, step 723, and if it is the existing table entry is replaced with a new entry, step 724. If it is not replacement of an old routine, a determination is then made if it is a new routine, step 725, and if it is, a new vector table is created with new entries, step 726. If it is not a replacement of an old routine or new functionality,

-30-

then it must be a new routine for an existing function. In this case, the new entries are simply added to the existing vector table, step 727.

Adding new entries to an existing vector table is accomplished by re-sizing the pointer to the vector table to make room for the new entries. In the currently preferred embodiment, in situations where the pointer cannot be re-sized, a new pointer is allocated and the old vector table is copied to the new location and then the new entries are added to the table and finally the low memory vector table pointer is updated with the new location. Other implementations may use other techniques, but would not depart from the spirit and scope of the present invention.

The patching technique described is used for each of the vector types. Use of other tables or additional indirection, e.g. a pointer to another table, would not depart from the spirit and scope of the present invention.

While the present invention has been described with reference to a computer operating system and Figures 1-7, it will be appreciated that the figures are for illustration only, and do not limit the spirit and scope of the present invention. In addition, it will be noted that the present invention may be realized using a variety of computer programming languages and hardware, and is not limited to any particular hardware and software configuration. The present invention may be utilized in any embodiment which has code stored in a read only storage device such as a ROM. For example, a microprocessor controller for controlling various operations of

-31-

an automobile may embody the present invention. Similar types of embodiments would be within the scope of the present invention.

Thus, a method for vectorizing object files for storage in a static storage device is disclosed.

-32-

CLAIMS**We Claim:**

1. A method for applying patches to code residing on a Read Only Memory (ROM) device, said code organized so that each reference to an entry point is made through a location in a Random Access Memory (RAM) device, said ROM device having code for one of a plurality of ROM release levels, said method comprising the steps of:
 - a) generating patch resources for each of said plurality of ROM release levels, each of said patch resources having a plurality of entries, each of said entries including addressing information for identifying a location in RAM through which said routine is accessed and a new routine;
 - b) identifying a ROM release for a ROM device to be patched;
 - c) retrieving a patch resource corresponding to said identified ROM release;
for each entry in said patch resource
 - d) updating an existing vector table entry if said new routine replaces an existing routine;
 - e) adding entries to a vector table if said new routine is added to an existing function; and
 - f) creating a new vector table if said new routine is for an added function.
2. The method as recited in Claim 1 wherein said step of generating patching resources for each of said plurality of ROM release levels is further comprised of the steps of:

-33-

- a) determining if a routine in said new ROM release level has a corresponding routine in said previous ROM release level;
- b) if a corresponding routine does exist, comparing the corresponding routines to determine if a difference exists, and if a difference does exist, creating a patch resource entry for said routine; and
- c) if a corresponding routine does not exist, creating a patch resource entry for said new routine based on addressing information from said new ROM release level.

3. The method as recited in Claim 2 wherein said addressing information for an entry in said patch resource is further comprised of a vector table pointer to a vector table and an offset into said vector table.

4. The method as recited in Claim 3 wherein said step of updating an existing vector table entry if said new routine replaces an existing routine is further comprised the of steps of:

- a) identifying a vector table from said vector table pointer of said patch resource entry;
- b) identifying a vector table entry from said offset of said patch resource entry;
- c) determining a load address for said new routine; and
- d) replacing the contents of said vector table entry with said load address.

-34-

5. The method as recited in Claim 4 wherein said step of adding entries to a vector table if said new routine is added to an existing function is further comprised of the steps of:

- a) identifying a vector table from said vector table pointer of said patch resource entry;
- b) determining a load address for said additional routine;
- c) adding an entry into said vector table; and
- d) inserting said load address into said added entry.

6. The method as recited in Claim 5 wherein said step of creating a new vector table if said new routine is for an added function is further comprised of the steps of:

- a) identifying a new vector table from said vector table pointer of said patch resource entry;
- b) allocating storage for said new vector table;
- c) determining a load address for said new routine;
- d) adding an entry into said new vector table; and
- e) inserting said load address into said added entry.

7. In an information processing system having a first version of operating software stored in a vectorized Read Only Memory (ROM) device, said operating software comprised of a plurality of routines which are accessed through a vector table, a method for starting a computer system so as to install patches for utilizing a second version of said operating software, said method comprising the steps of:

-35-

a) retrieving a patch resource , said patch resource comprised of a plurality of entries for replacing routines, adding new routines for existing functions or adding new routines for new function, each of said entries comprised of address information for identifying an entry in a vector table and a routine;

for each entry in said patch resource:

b) updating a vector table entry if said patch resource entry is for replacing a routine;

c) adding entries to a vector table if said patch resource entry is for adding new routines to existing functions; and

d) creating a new vector table if said patch resource entry is for adding new routines for new function.

8. The method as recited in Claim 7 wherein said step of updating a vector table entry if said patch resource entry is for replacing a routine is further comprised of the steps of:

a) determining a first load address for a routine in said patch resource entry;

b) determining a vector table entry location from said address information of said vector patch resource entry; and

c) entering said first load address into said vector table entry location.

-36-

9. The method as recited in Claim 7 wherein said step of adding entries to a vector table if said patch resource entry is for adding new routines is further comprised of the steps of:

- a) determining a second load address for a routine in said patch resource entry;
- b) creating a new vector table entry from said address information of said vector patch resource entry; and
- c) entering said second load address into said vector table entry location.

10. The method as recited in Claim 7 wherein said step of creating a new vector table if said patch resource entry is for adding new function is further comprised of the steps of:

- a) determining a new vector table address from said address information of said vector patch resource entry;
- b) creating said new vector table;
- c) determining a third load address for said routine in said vector patch resource entry;
- d) adding an entry into said new vector table; and
- e) inserting said third load address into said added entry.

11. For a family of computer systems, each computer system in said family of computer systems having a vectorized Read Only Memory (ROM) device for storing operating instructions, said operating instructions comprised of a plurality of routines, each of said routines being accessed

-37-

through a vector table, a method for generating patching resources for said operating software, said patching resources for updating said operating software to a desired version level, said method comprising the steps of:

a) assigning a ROM version number for each ROM device, said ROM version number corresponding to the version of the operating instructions residing on said ROM;

b) assigning X as a version number for said new version of said ROM;

for a new version of the ROM performing the steps of:

c) comparing version X of the ROM with previous ROM versions;

d) for each previous ROM version, generating a patch resource based on the differences with said previous ROM versions, said patching resource comprised of a plurality of entries, each of said entries comprising vector table addressing information and a routine.

12. The method as recited in Claim 11 wherein said step of for each previous ROM version, generating a patch resource based on the differences with said previous ROM versions is further comprised of the step of:

a) determining if a routine in said new ROM release level has a corresponding routine in said previous ROM release level;

b) if a corresponding routine does exist, comparing the corresponding routines to determine if a difference exists, and if a difference exists, creating a patch resource entry for said routine; and

-38-

c) if a corresponding routine does not exist, creating a patch resource entry for said new routine based on addressing information from said new ROM release level.

13. The method as recited in Claim 12 wherein said addressing information is a vector table pointer for identifying a location for said vector table and a vector table offset for identifying an entry in said vector table.

14. The method as recited in Claim 13 wherein said step of creating a patch resource entry for said new routine based on addressing information from said new ROM release level is further comprised of retrieving a vector table pointer and a vector table offset from said new ROM release level.

SUBSTITUTE SHEET (RULE 26)

1/8

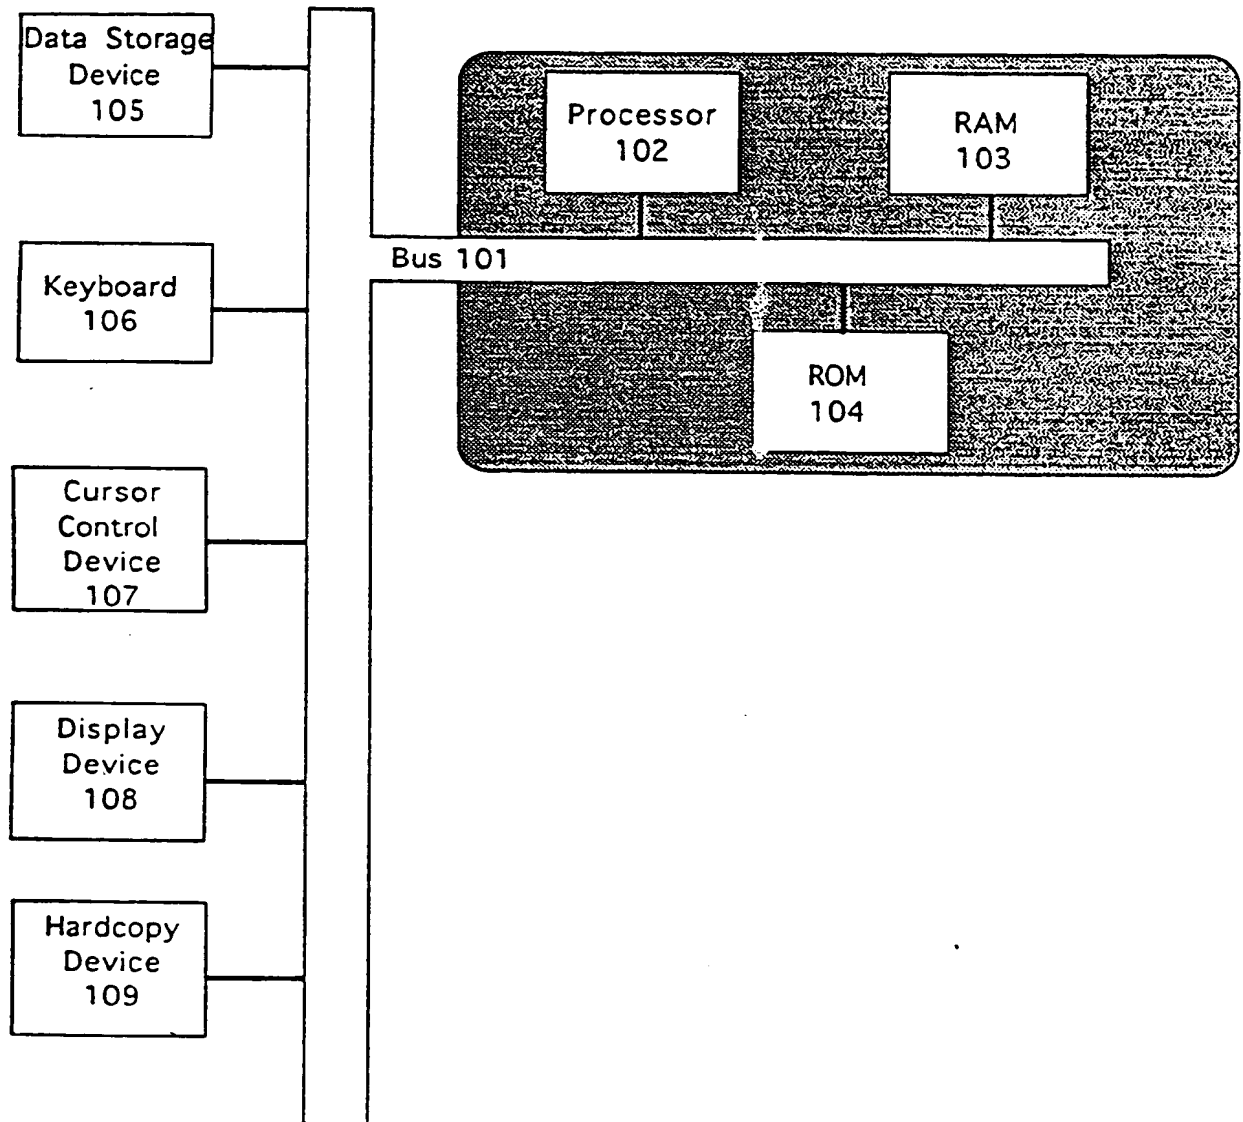




Figure 1

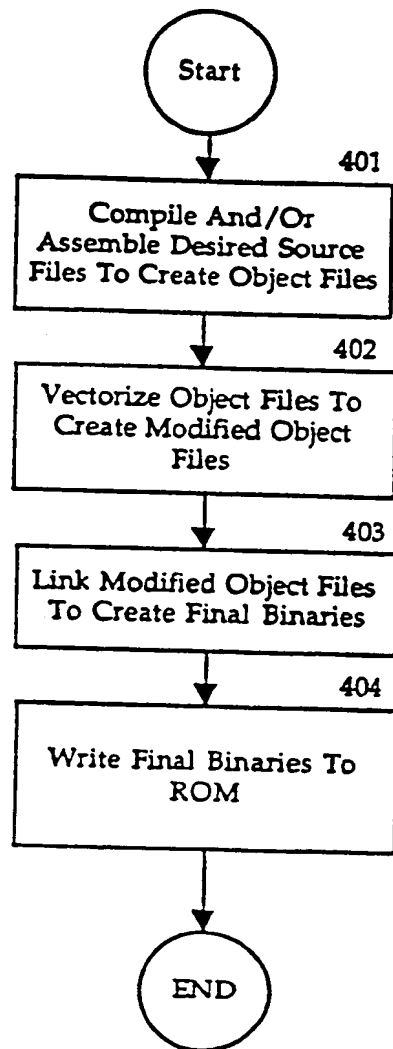
Figure 2

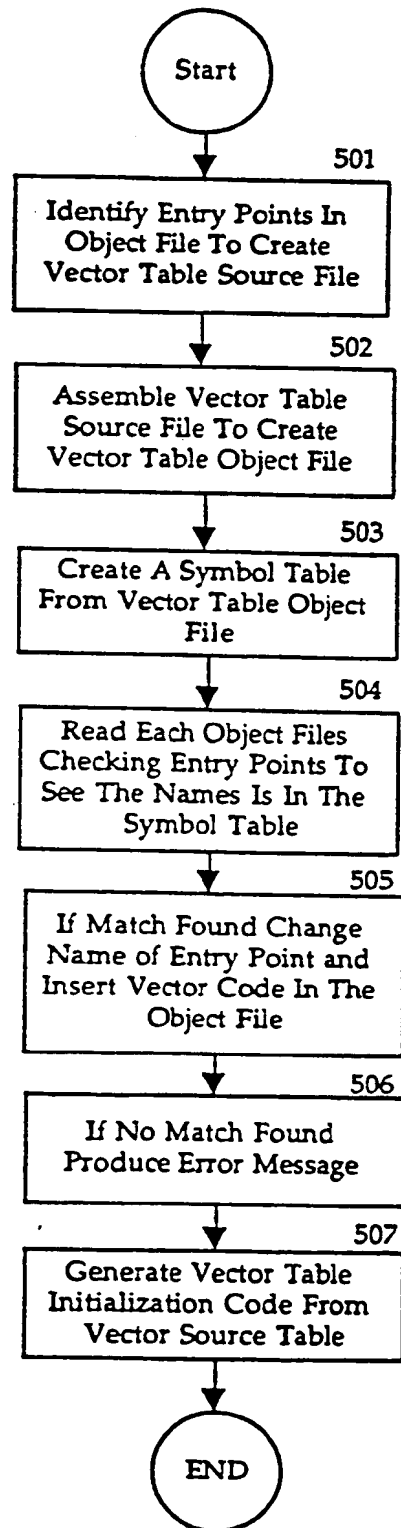




4/8


Figure 4



5/8
Figure 5

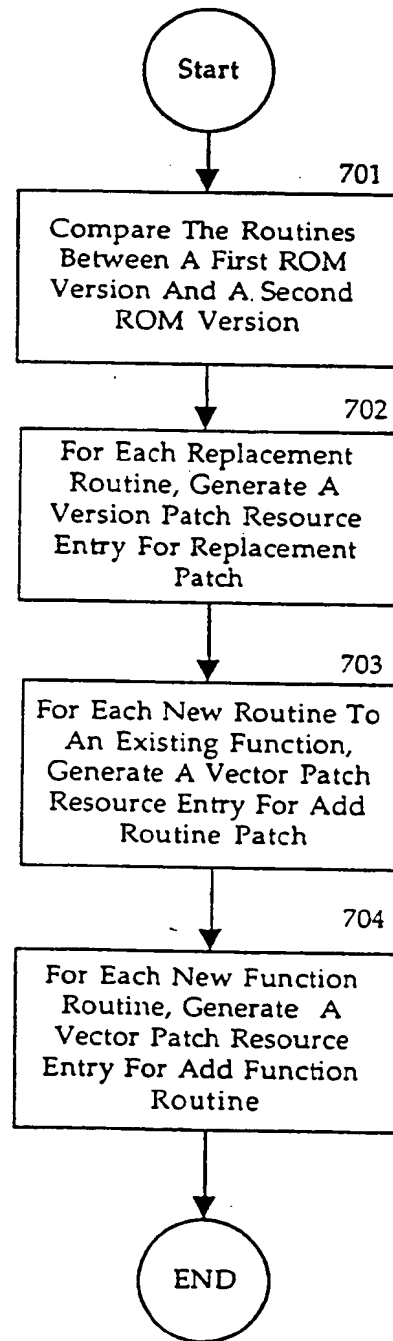
6/8

Figure 6



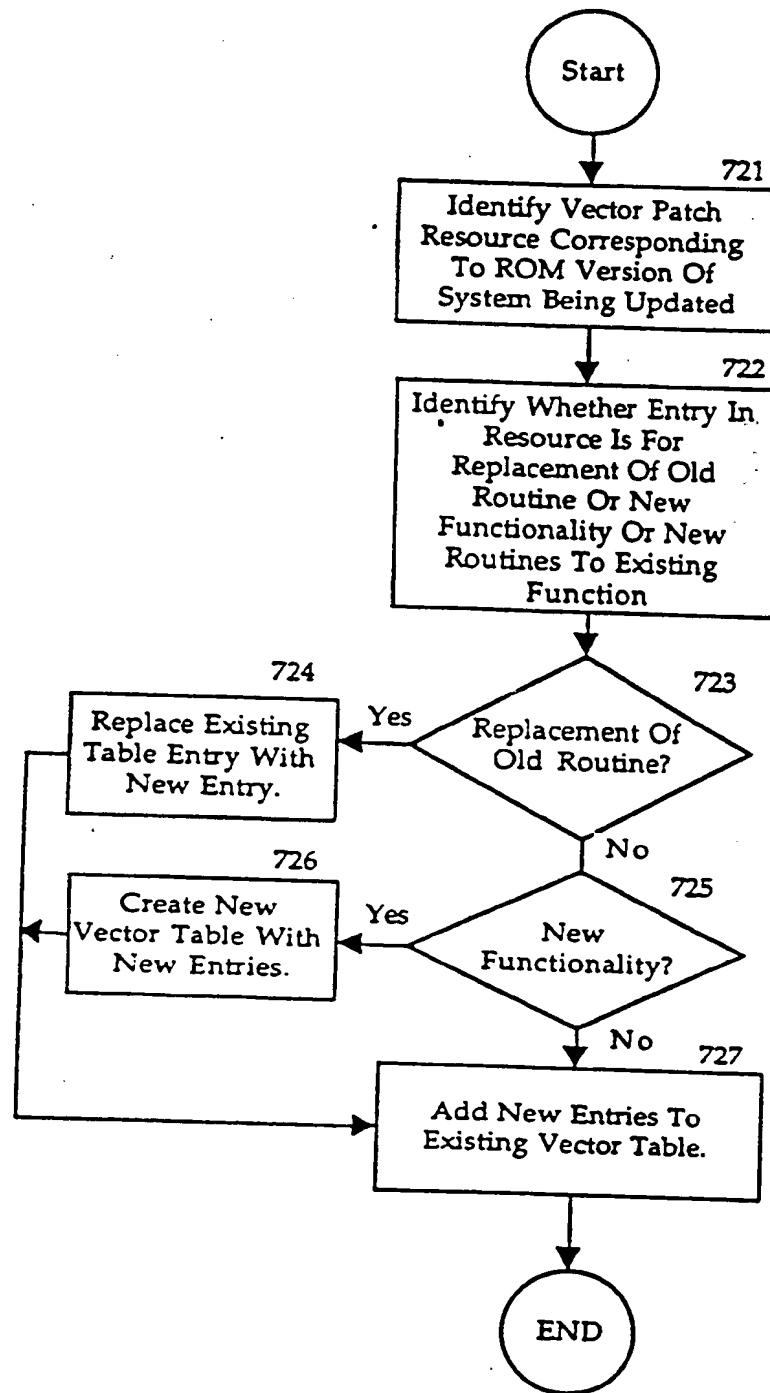
7/8

Figure 7a



8/8

Figure 7b



INTERNATIONAL SEARCH REPORT

Internat. Application No
PCT/US 94/04995

A. CLASSIFICATION OF SUBJECT MATTER
IPC 5 G06F9/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 5 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	IBM TECHNICAL DISCLOSURE BULLETIN. vol. 31, no. 1, June 1988, NEW YORK US pages 294 - 298 'dual indirect ram/rom jump tables for firmware updates' see the whole document -----	1,7,11
A	EP,A,0 321 000 (NEC CORP.) 21 June 1989 see the whole document -----	1,3-10
A	EP,A,0 384 357 (OMROM CORP.) 29 August 1990 -----	1,7
A	EP,A,0 472 812 (LANDIS & GYR BETREIBS AG) 4 March 1992 * page 3, lines 39-49; page 5, lines 21-42 * -----	11

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- *&* document member of the same patent family

Date of the actual completion of the international search

29 August 1994

Date of mailing of the international search report

07.09.94

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+ 31-70) 340-3016

Authorized officer:

Weinberg, L

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 94/04995

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP-A-0321000	21-06-89	JP-A- 1161534 JP-A- 1177128 US-A- 5182807	26-06-89 13-07-89 26-01-93
EP-A-0384357	29-08-90	JP-A- 2220158	03-09-90
EP-A-0472812	04-03-92	NONE	